arXiv:1701.02208v1 [math.AT] 9 Jan 2017

# Barcodes of Towers and a Streaming Algorithm for Persistent Homology

Michael Kerber * Hannah Schreiber *

January 10, 2017

**Abstract**

A tower is a sequence of simplicial complexes connected by simplicial maps. We show how to compute a filtration, a sequence of nested simplicial complexes, with the same persistent barcode as the tower. Our approach is based on the coning strategy by Dey et al. (SoCG 2014). We show that a variant of this approach yields a filtration that is asymptotically only marginally larger than the tower and can be efficiently computed by a streaming algorithm, both in theory and in practice. Furthermore, we show that our approach can be combined with a streaming algorithm to compute the barcode of the tower via matrix reduction. The space complexity of the algorithm does not depend on the length of the tower, but the maximal size of any subcomplex within the tower. Experimental evaluations show that our approach can efficiently handle towers with billions of complexes.

## 1 Introduction

**Motivation and problem statement.** Persistent homology [17, 6, 16] is a paradigm to analyze how topological properties of general data sets evolve across multiple scales. Thanks to the success of the theory in finding applications in many sciences (see, e.g., [28, 21] for recent enumerations), there is a growing demand for efficient computations of the involved topological invariants.

In this paper, we consider a sequence of simplicial complexes $(\mathbb{K}_i)_{i=0,\dots,m}$ and simplicial maps $\phi_i : \mathbb{K}_i \to \mathbb{K}_{i+1}$ connecting them, calling this data a *(simplicial) tower* of length $m$. Applying the homology functor with an arbitrary field, we obtain a *persistence module*, a sequence of vector spaces connected by linear maps. Such a module decomposes into a *barcode*, a collection of intervals, each representing a homological feature in the tower that spans over the specified range of scales.

Our computational problem is to compute the barcode of a given tower efficiently. The most prominent case of a tower is when all maps $f_i$ are inclusion maps. In this case one obtains a *filtration*, a sequence of nested simplicial complexes. A considerable amount of work went into the study of fast algorithms for the filtration case, which culminated in fast software libraries for this task. The more general case of towers recently received growing interest in the context of sparsification technique for the Vietoris-Rips and Čech complexes; see the related work section below for a detailed discussion.

**Results.** As our first result, we show that any tower can be *efficiently* converted into a *small* filtration with the same barcode. Using the well-known concept of *mapping cylinders* from algebraic topology [20], it is easy to see that such a conversion is possible in principle. Dey, Fan, and Wang [14] give an explicit construction, called "coning", for the generalized case of *zigzag towers*. Using a simple variant of their coning strategy, we obtain a filtration whose size is only marginally larger than the length of the tower in the worst case. Furthermore, we experimentally show that the size is even smaller on realistic instances.

*Graz University of Technology, Austria

To describe our improved coning strategy, we discuss the case that a simplicial map in the tower contracts two vertices $u$ and $v$. The coning strategy by Dey et al. proposes to join $u$ with the closed star of $v$, making all incident simplices of $v$ incident to $u$ without changing the homotopy type. The vertex $u$ is then taken as the representative of the contracted pair in the further processing of the tower. We refer to the number of simplices that the join operation adds to the complex as the *cost* of the contraction. Quite obviously, the method is symmetric in $u$ and $v$, and we have two choices to pick the representative, leading to potentially quite different costs. We employ the self-evident strategy to pick the representative that leads to smaller costs (somewhat reminiscent of the "union-by-weight" rule in the union-find data structure [12, §21]). Perhaps surprisingly, this idea leads to an asymptotically improved size bound on the filtration. We prove this by an abstraction to path decompositions on weighted forest which might be of some independent interest. Altogether, the worst-case size of the filtration is $O(\Delta \cdot n \cdot \log(n_0))$, where $\Delta$ is the maximal dimension of any complex in the tower, and $n/n_0$ is the number of simplices/vertices added to the tower.

We also provide a conversion algorithm whose time complexity is roughly proportional to the total number of simplices in the resulting filtration. One immediate benefit is a generic solution to compute barcodes of towers: just convert the tower to a filtration and apply one of the efficient implementations for barcodes of filtrations. Indeed, we experimentally show that on not-too-large towers, our approach is competitive with, and sometimes outperforms Simpers, an alternative approach that computes the barcode of towers with *annotations*, a variant of the persistent cohomology algorithm.

Our second contribution is an space-efficient version of the just mentioned algorithmic pipeline that is applicable to very large towers. To motivate the result, let the *width* of a tower denote the maximal size of any simplicial complex among the $\mathbb{K}_i$. Consider a tower with a very large length (say, larger than the number of bytes in main memory) whose width remains relatively small. In this case, our conversion algorithm yields a filtration that is very large as well. Most existing implementations for barcode computation read the entire filtration on initialization and must be converted to streaming algorithm to handle such instances. Moreover, algorithms based on matrix reduction are required to keep previously reduced columns because they might be needed in subsequent reduction steps. This leads to a high memory consumption for the barcode computation.

We show that with minor modifications, the standard persistent algorithm can be turned into a streaming algorithm with smaller space complexity in the case of towers. The idea is that upon contractions, simplices become *inactive* and cannot get additional cofaces. Our approach makes use of this observation by modifying the boundary matrix such that columns associated to inactive simplices can be removed. Combined with our conversion algorithm, we can compute the barcode of a tower of width $\omega$ keeping only up to $O(\omega)$ columns of the boundary matrix in memory. This yields a space complexity of $O(\omega^2)$ and a time complexity of $O((\Delta \cdot n \cdot \log(n_0))\omega^2)$ in the worst case. We implemented a practically improved variant that makes use of additional heuristics to speed up the barcode computation in practice and resembles the *chunk algorithm* presented in [1].

We tested our implementation on various challenging data sets. The source code of the implementation is available at `https://bitbucket.org/schreiberh/sophia/`.

**Related work.** Already the first works on persistent homology pointed out the existence of efficient algorithm to compute the barcode invariant (or equivalently, the persistent diagram) for filtrations [17, 30]. As a variant of Gaussian elimination, the worst-case complexity is cubic. Remarkable theoretical follow-up results are a persistence algorithm in matrix multiplication time [25], an output-sensitive algorithm to compute only high-persistent features with linear space complexity [10], and a conditional lower bound on the complexity relating the problem to rank computations of sparse matrices [18].

On realistic instances, the standard algorithm has shown a quasi-linear behavior in practice despite its pessimistic worst-case complexity. Nevertheless, many improvements of the standard algorithm have been presented in the last years which improve the runtime by several orders of magnitude. One line of research exploits the special structure of the boundary matrix to speed up the reduction process [9]. This idea has led to efficient parallel algorithms for persistence in shared [1] and distributed memory [2]. Moreover, of same importance as the reduction strategy is an appropriate choice of data structures in the reduction process as demonstrated by the Phat library [3]. A parallel development was the development of dual algorithms using persistent

cohomology, based on the observation that the resulting barcode is identical [13]. The *annotation algorithm* [14, 4] is an optimized variant of this idea realized in the GUDHI library [23]. It is commonly considered as an advantage of annotations that only a cohomology basis must be kept during the reduction process, making it more space efficient than reduction-based approaches. We refer to the comparative study [27] for further approaches and software for persistence on filtrations.

Moreover, generalizations of the persistence paradigm are an active field of study. *Zigzag persistence* is a variant of persistence where the maps in the filtration are allowed to map in either direction (that is, either $\phi_i : \mathbb{K}_i \hookrightarrow \mathbb{K}_{i+1}$ or $\phi_i : \mathbb{K}_i \hookleftarrow \mathbb{K}_{i+1}$). The barcode of zigzag filtrations is well-defined [7] as a consequence of Gabriel's theorem [19] on decomposable quivers – see [28] for a comprehensive introduction. The initial algorithms to compute this barcode [8] has been improved recently [24]. Our case of towers of complexes and simplicial maps can be modeled as a zigzag filtration and therefore sits in-between the standard and the zigzag filtration case.

Dey et al. [14] described the first efficient algorithm to compute the barcode of towers. Instead of the aforementioned coning approach explained in their paper, their implementation handles contractions with an empirically smaller number of insertions, based on the link condition. Recently, the authors have released the SIMPERS library[1] that implements their annotation algorithm from the paper.

The case of towers has received recent attention in the context of approximate Vietoris-Rips and Čech filtrations. The motivation for approximation is that the (exact) topological analysis of a set of $n$ points in $d$-dimensions requires a filtration of size $O(n^{d+1})$ which is prohibitive for most interesting input sizes. Instead, one aims for a filtration or tower of much smaller size, with the guarantee that the approximate barcode will be close to the exact barcode ("close" usually means that the bottleneck distance between the barcodes on the logarithmic scale is upper bounded; we refer to the cited works for details). The first such type of result by Sheehy [29] resulted in a approximate filtration; however, it has been observed that passing to towers allows more freedom in defining the approximation complexes and somewhat simplifies the approximation schemes conceptually. See [14, 5, 22, 11] for examples. Very recently, the SimBa library [15] brings these theoretical approximation techniques for Vietoris-Rips complexes into practice. The approach consists of a geometric layer to compute a tower, and an algebraic layer to compute its barcode, for which they use SIMPERS. Our approach can be seen as an alternative realization of this algebraic layer.

**Outline.** We introduce the necessary basic concepts in Section 2. We describe our conversion algorithm from general towers to barcodes in Section 3. The streaming algorithm for persistence is discussed in Section 4.

# 2 Background

**Simplicial Complexes.** Given a finite *vertex set* $V$, a *simplex* is merely a non-empty subset of $V$; more precisely, a *$k$-simplex* is a subset consisting of $k+1$ vertices, and $k$ is called the *dimension* of the simplex. Throughout the paper, we will denote simplices by small Greek letters, except for vertices (0-simplices) which we denote by $u$, $v$, and $w$. For a $k$-simplex $\sigma$, a simplex $\tau$ is a *face* of $\sigma$ if $\tau \subseteq \sigma$. If $\tau$ is of dimension $\ell$, we call it a $\ell$-face. If $\ell < k$, we call $\tau$ a *proper face* of $\sigma$, and if $\ell = k - 1$, we call it a *facet*. For a simplex $\sigma$ and a vertex $v \notin \sigma$, we define the *join* $v * \sigma$ as the simplex $\{v\} \cup \sigma$. These definition are inspired by visualizing a $k$-simplex as the convex hull of $k + 1$ affinely independent points in $\mathbb{R}^k$, but we will not need this geometric interpretation in our arguments.

An *(abstract) simplicial complex* $\mathbb{K}$ over $V$ is a set of simplices that is closed under taking faces. We call $V$ the *vertex set* of $\mathbb{K}$ and write $\mathcal{V}(\mathbb{K}) := V$. The *dimension* of $\mathbb{K}$ is the maximal dimension of its simplices. For $\sigma, \tau \in \mathbb{K}$, we call $\sigma$ a *coface* of $\tau$ in $\mathbb{K}$ if $\tau$ is a face of $\sigma$. In this case, $\sigma$ is a *cofacet* of $\tau$ if their dimensions differ by exactly one. A simplicial complex $\mathbb{L}$ is a *subcomplex* of $\mathbb{K}$ if $\mathbb{L} \subseteq \mathbb{K}$. Given $\mathcal{W} \subseteq \mathcal{V}$, the *induced subcomplex* by $\mathcal{W}$ is the set of all simplices $\sigma$ in $\mathbb{K}$ with $\sigma \subseteq \mathcal{W}$. For a subcomplex $\mathbb{L} \subseteq \mathbb{K}$ and a vertex $v \in \mathcal{V}(\mathbb{K}) \setminus \mathcal{V}(\mathbb{L})$, we define the join $v * \mathbb{L} := \{v * \sigma \mid \sigma \in \mathbb{L}\}$. For a vertex $v \in \mathbb{K}$, the *star* of $v$ in $\mathbb{K}$, denoted by $\mathrm{St}(v, \mathbb{K})$, is the set of all cofaces of $v$ in $\mathbb{K}$. In general, the star is not a subcomplex, but we can make it a subcomplex by adding all faces of star

[1] `http://web.cse.ohio-state.edu/~tamaldey/SimPers/Simpers.html`

simplices, which is denoted by the *closed star* $\overline{\mathrm{St}}(v, \mathbb{K})$. Equivalently, the closed star is the smallest subcomplex of $\mathbb{K}$ containing the star of $v$. The *link* of $v$, $\mathrm{Lk}(v, \mathbb{K})$, is defined as $\overline{\mathrm{St}}(v, \mathbb{K}) \setminus \mathrm{St}(v, \mathbb{K})$. It can be checked that the link is a subcomplex of $\mathbb{K}$. When the complex is clear from context, we will sometimes omit the $\mathbb{K}$ in the notation of stars and links.

**Simplicial maps.** A map $\mathbb{K} \xrightarrow{\phi} \mathbb{L}$ between simplicial complexes is called *simplicial* if with $\sigma = \{v_0, \ldots, v_k\} \in \mathbb{K}$, $\phi(\sigma)$ is equal to $\{\phi(v_0), \ldots, \phi(v_k)\}$ and $\phi(\sigma)$ is a simplex in $\mathbb{L}$. By definition, a simplicial map maps vertices to vertices and is completely determined by its action on the vertices. Moreover, the composition of simplicial maps is again simplicial.

A simple example of a simplicial map is the inclusion map $\mathbb{L} \xhookrightarrow{\phi} \mathbb{K}$ where $\mathbb{L}$ is a subcomplex of $\mathbb{K}$. If $\mathbb{K} = \mathbb{L} \cup \{\sigma\}$ with $\sigma \notin \mathbb{L}$, we call $\phi$ an *elementary inclusion*. The simplest example of a non-inclusion simplicial map is $\mathbb{K} \xrightarrow{\phi} \mathbb{L}$ such that there exist two vertices $u, v \in \mathbb{K}$ with $\mathcal{V}(\mathbb{L}) = \mathcal{V}(\mathbb{K}) \setminus \{v\}$, $\phi(u) = \phi(v) = u$, and $\phi$ is the identity on all remaining vertices of $\mathbb{K}$. We call call $\phi$ an *elementary contraction* and write $(u, v) \rightsquigarrow u$ as a shortcut. It is well-known that any simplicial map $\mathbb{K} \xrightarrow{\phi} \mathbb{L}$ can be written as the composition of elementary contractions and elementary inclusions.

A *tower* of length $m$ is a collection of simplicial complexes $\mathbb{K}_0, \ldots, \mathbb{K}_m$ and simplicial maps $\phi_i : \mathbb{K}_i \to \mathbb{K}_{i+1}$ for $i = 0, \ldots, m-1$. From this initial data, we obtain simplicial maps $\phi_{i,j} : \mathbb{K}_i \to \mathbb{K}_j$ for $i \leq j$ by composition, where $\phi_{i,i}$ is simply the identity map on $\mathbb{K}_i$. The term "tower" is taken from category theory, where it denotes a (directed) path in a category with morphisms from objects with smaller indices to objects with larger indices. Indeed, since simplicial complexes form a category with simplicial maps as their morphisms, the specified data defines a tower in this category. A tower is called a *filtration* if all $\phi_i$ are inclusion maps. The *dimension* of a tower is the maximal dimension among the $\mathbb{K}_i$, and the *width* of a tower is the maximal size among the $\mathbb{K}_i$. For filtrations, dimension and width are determined by the dimension and size of $\mathbb{K}_m$, but this is not necessarily true for general towers.

**Homology and Collapses.** For a fixed base field $\mathbb{F}$, let $H_p(\mathbb{K}) := H_p(\mathbb{K}, \mathbb{F})$ the *p-dimensional homology group* of $\mathbb{K}$. It is well-known that $H_p(\mathbb{K})$ is a $\mathbb{F}$-vector space. Moreover, a simplicial map $\mathbb{K} \xrightarrow{\phi} \mathbb{L}$ induces a linear map $H_p(\mathbb{K}) \xrightarrow{\phi^*} H_p(\mathbb{L})$. In categorical terms, the equivalent statement is that homology is a functor from the category of simplicial complexes and simplicial maps to the category of vector spaces and linear maps.

We will make use of the following homology-preserving operation: a *free face* in $\mathbb{K}$, is a simplex with exactly one proper coface in $\mathbb{K}$. An *elementary collapse* in $\mathbb{K}$ is the operation of removing a free face and its unique coface from $\mathbb{K}$, yielding a subcomplex of $\mathbb{K}$. We say that $\mathbb{K}$ *collapses to* $\mathbb{L}$, if there is a sequence of elementary collapses transforming $\mathbb{K}$ into $\mathbb{L}$. The following result is then well-known:

**Lemma 1.** *Let $\mathbb{K}$ be a complex that collapses into the complex $\mathbb{L}$. Then, the inclusion map $\mathbb{L} \xhookrightarrow{\phi} \mathbb{K}$ induces an isomorphism $\phi_*$ between $H_p(\mathbb{L})$ and $H_p(\mathbb{K})$.*

**Barcodes.** A *persistence module* is a sequence vector spaces $\mathbb{V}_0, \ldots, \mathbb{V}_m$ and linear maps $f_{i,j} : \mathbb{V}_i \to \mathbb{V}_j$ for $i < j$ such that $f_{i,i} = \mathrm{id}_{\mathbb{V}_i}$ and $f_{i,k} = f_{j,k} \circ f_{i,j}$ for $i \leq k \leq j$. As primary example, we obtain a persistence module by applying the homology functor on any simplicial tower. Persistence modules admit a decomposition into indecomposable summands in the following sense. Writing $I_{b,d}$ with $b \leq d$ for the persistence module

$$\underbrace{0 \xrightarrow{0} \ldots \xrightarrow{0} 0}_{b-1 \text{ times}} \xrightarrow{0} \underbrace{\mathbb{F} \xrightarrow{\mathrm{id}} \ldots \xrightarrow{\mathrm{id}} \mathbb{F}}_{d-b+1 \text{ times}} \xrightarrow{0} \underbrace{0 \xrightarrow{0} \ldots \xrightarrow{0} 0}_{m-d \text{ times}},$$

we can write every persistence module as the direct sum $I_{b_1,d_1} \oplus \ldots \oplus I_{b_s,d_s}$, where the direct sum of persistence modules is defined component-wise for vector spaces and linear maps in the obvious way. Moreover, this decomposition is uniquely defined up to isomorphisms and re-ordering, thus the pairs $(b_1, d_1), \ldots, (b_s, d_s)$ are an invariant of the persistence module, called its *barcode*. When the persistence module was generated by a tower, we also talk about the barcode of the tower.

**Matrix reduction.** In this paragraph, we assume that $(\mathbb{K}_i)_{i=0,\dots,m}$ is a filtration such that $\mathbb{K}_0 = \emptyset$ and $\mathbb{K}_{i+1}$ has exactly one more simplex than $\mathbb{K}_i$. We label the simplices of $\mathbb{K}_m$ accordingly as $\sigma_1, \dots, \sigma_m$, with $\mathbb{K}_i \setminus \mathbb{K}_{i-1} = \{\sigma_i\}$. The filtration can be encoded as a *boundary matrix* $\partial$ of dimension $m \times m$, where the $(ij)$-entry is 1 if $\sigma_i$ is a facet of $\sigma_j$, and 0 otherwise. In other words, the $j$-th column of $\partial$ encodes the facets of $\sigma_j$, and the $i$-th row of $\partial$ encodes the cofacets of $\sigma_i$. Moreover, $\partial$ is upper-triangular because every $\mathbb{K}_i$ is a simplicial complex. We will sometimes identify rows and columns in $\partial$ with the corresponding simplex in $\mathbb{K}_m$. Adding the $k$-simplex $\sigma_i$ to $\mathbb{K}_{i-1}$ either introduces one new homology class (of dimension $k$), or turns a non-trivial homology class (of dimension $k-1$) trivial. We call $\sigma_i$ and the $i$-th column of $\partial$ *positive* or *negative*, respectively (with respect to the given filtration).

For the computation of the barcode, we assume for simplicity homology over the base field $\mathbb{Z}_2$, and interpret the coefficients of $\partial$ accordingly. In an arbitrary matrix $A$, a *left-to-right column addition* is an operation of the form $A_k \leftarrow A_k + A_\ell$ with $\ell < k$, where $A_k$ and $A_\ell$ are columns of the matrix. The *pivot* of a non-zero column is the largest non-zero index of the corresponding column. A non-zero entry is called a pivot if its row is the pivot of the column. A matrix $R$ is called a *reduction* of $A$ if $R$ is obtained by a sequence of left-to-right column additions from $A$ and no two columns in $R$ have the same pivot. It is well-known that, although $\partial$ does not have a unique reduction, the pivots of all its reductions are the same. Moreover, the pivots $(b_1, d_1), \dots, (b_s, d_s)$ of $R$ are precisely the barcode of the filtration. A direct consequence is that a simplex $\sigma_i$ is positive if and only if the $i$-th column in $R$ is zero.

The standard persistence algorithm processes the columns from left to right. In the $j$-th iteration, as long as the $j$-th column is not empty and has a pivot that appears in a previous column, it performs a left-to-right column addition. In this work, we use a simple improvement of this algorithm that is called *compression*: before reducing the $j$-th column, it first scans through the non-zero entries of the column. If a row index $i$ corresponds to a negative simplex (i.e., if the $i$-th column is not zero at this point in the algorithm), the row index can be deleted without changing the pivots of the matrix. After this initial scan, the column is reduced in the same way as in the standard algorithm. See [1, §. 3] for a discussion (we remark that this optimization was also used in [30]).

# 3 From towers to filtrations

We phrase now our first result which says that any tower can be converted into a filtration of only marginally larger size with a space-efficient streaming algorithms:

**Theorem 2** (Conversion Theorem)**.** *Let*

$$\mathcal{T}: \quad \mathbb{K}_0 \xrightarrow{\phi_0} \mathbb{K}_1 \xrightarrow{\phi_1} \dots \xrightarrow{\phi_{m-1}} \mathbb{K}_m$$

*be a tower where, w.l.o.g., $\mathbb{K}_0 = \emptyset$ and each $\phi_i$ is either an elementary inclusion or an elementary contraction. Let $\Delta$ denote the dimension and $\omega$ the width of the tower, and let $n \leq m$ denote the total number of elementary inclusions, and $n_0$ the number of vertex inclusions. Then, there exists a filtration*

$$\mathcal{F}: \quad \hat{\mathbb{K}}_0 \hookrightarrow \hat{\mathbb{K}}_1 \hookrightarrow \dots \hookrightarrow \hat{\mathbb{K}}_m$$

*where the inclusions are not necessarily elementary, such that $\mathcal{T}$ and $\mathcal{F}$ have the same barcode and the width of the filtration $|\hat{\mathbb{K}}_m|$ is at most $O(\Delta \cdot n \log n_0)$. Moreover, $\mathcal{F}$ can be computed from $\mathcal{T}$ with a streaming algorithm in $O(\Delta \cdot |\hat{\mathbb{K}}_m| \cdot C_\omega)$ time and space complexity $O(\Delta \cdot \omega)$, where $C_\omega$ is the cost of an operation in a dictionary with $\omega$ elements.*

The remainder of the section is organized as follows. We define $\mathcal{F}$ in Section 3.1 and prove that it yields the same barcode as $\mathcal{T}$ in Section 3.2. In Section 3.3, we prove the upper bound on the width of the filtration. In Section 3.4, we explain the algorithm to compute $\mathcal{F}$ and analyze its time and space complexity.

### 3.1 Active and small coning

**Coning.** We briefly revisit the *coning strategy* introduced by Dey, Fan and Wang [14]. Let $\phi : \mathbb{K} \to \mathbb{L}$ be an elementary contraction $(u, v) \rightsquigarrow u$ and define

$$\mathbb{L}^* = \mathbb{K} \cup \left(u * \overline{\mathrm{St}}(v, \mathbb{K})\right).$$

An example is shown in Figure 1.

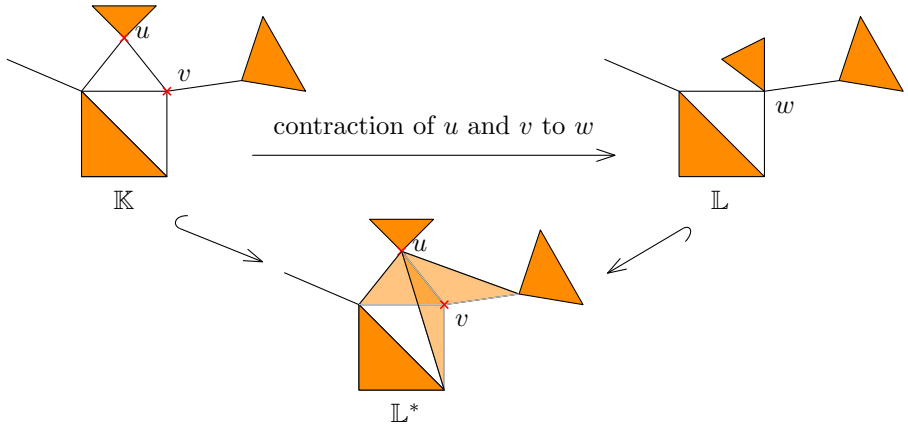


Figure 1: Construction example of $\mathbb{L}^*$, were $u$ and $v$ in $\mathbb{K}$ are contracted to $w$ in $\mathbb{L}$.

Dey et al. show that $\mathbb{L} \subseteq \mathbb{L}^*$ and that the map induced by inclusion is an isomorphism between $H(\mathbb{L})$ and $H(\mathbb{L}^*)$. By applying this result at any elementary contraction, this implies that every zigzag tower can be transformed into a zigzag filtration with identical barcode.

Given a tower $\mathcal{T}$, we can also obtain an non-zigzag filtration using coning, if we continue the operation on $\mathbb{L}^*$ instead of going back to $\mathbb{L}$. More precisely, we set $\tilde{\mathbb{K}}_0 := \emptyset$ and if $\phi_i$ is an inclusion of simplex $\sigma$, we set $\tilde{\mathbb{K}}_{i+1} := \tilde{\mathbb{K}}_i \cup \{\sigma\}$. If $\phi_i$ is a contraction $(u, v) \rightsquigarrow u$, we set $\tilde{\mathbb{K}}_{i+1} = \tilde{\mathbb{K}}_i \cup \left(u * \overline{\mathrm{St}}(v, \tilde{\mathbb{K}}_i)\right)$. Indeed, it can be proved that $(\tilde{\mathbb{K}}_i)_{i=0,\dots,m}$ has the same barcode as $\mathcal{T}$. However, the filtration will not be small, and we will define a smaller variant now.

Our new construction yields a sequence of complexes $\hat{\mathbb{K}}_0, \dots, \hat{\mathbb{K}}_m$ with $\hat{\mathbb{K}}_i \subseteq \hat{\mathbb{K}}_{i+1}$. During the construction, we maintain a flag for each vertex in $\hat{\mathbb{K}}_i$, which marks the vertex as *active* or *inactive*. A simplex is called *active* if all its vertices are active, and *inactive* otherwise. For a vertex $u$ and a complex $\hat{\mathbb{K}}_i$, let $\mathrm{Act}\overline{\mathrm{St}}(u, \hat{\mathbb{K}}_i)$ denote its *active closed star*, which is the set of active simplices in $\hat{\mathbb{K}}_i$ in the closed star of $u$.

The construction is inductive, starting with $\hat{\mathbb{K}}_0 := \emptyset$. If $\mathbb{K}_i \xrightarrow{\phi_i} \mathbb{K}_{i+1}$ is an elementary inclusion with $\mathbb{K}_{i+1} = \mathbb{K}_i \cup \{\sigma\}$, set $\hat{\mathbb{K}}_{i+1} := \hat{\mathbb{K}}_i \cup \{\sigma\}$. If $\sigma$ is a vertex, we mark it as active. It remains the case that $\mathbb{K}_i \xrightarrow{\phi_i} \mathbb{K}_{i+1}$ is an elementary contraction of the vertices $u$ and $v$. If $|\mathrm{Act}\overline{\mathrm{St}}(u, \hat{\mathbb{K}}_i)| \leq |\mathrm{Act}\overline{\mathrm{St}}(v, \hat{\mathbb{K}}_i)|$, we set

$$\hat{\mathbb{K}}_{i+1} = \hat{\mathbb{K}}_i \cup \left(v * \mathrm{Act}\overline{\mathrm{St}}(u, \hat{\mathbb{K}}_i)\right)$$

and mark $u$ as inactive. Otherwise, we set

$$\hat{\mathbb{K}}_{i+1} = \hat{\mathbb{K}}_i \cup \left(u * \mathrm{Act}\overline{\mathrm{St}}(v, \hat{\mathbb{K}}_i)\right)$$

and mark $v$ as inactive. This ends the description of the construction. We write $\mathcal{F}$ for the filtration $(\hat{\mathbb{K}}_i)_{i=0,\dots,m}$.

There are two major changes compared to the construction of $(\tilde{\mathbb{K}}_i)_{i=0,\dots,m}$. First, to counteract the potentially large growth of the involved cones, we restrict coning to active simplices. We will show below that the subcomplex of $\hat{\mathbb{K}}_i$ induced by the active vertices is isomorphic to $\mathbb{K}_i$. As a consequence, we add the same number of simplices by passing from $\hat{\mathbb{K}}_i$ to $\hat{\mathbb{K}}_{i+1}$ as in the approach by Dey et al. does when passing from $\mathbb{K}$ to $\mathbb{L}^*$.

A second difference is that our strategy exploits that an elementary contraction of two vertices $u$ and $v$ leaves us with a choice: we can either take $u$ or $v$ as the representative of the contracted

vertex. In terms of simplicial maps, these two choices correspond to setting $\phi_i(u) = \phi_i(v) = u$ or $\phi_i(u) = \phi_i(v) = v$, if $\phi_i$ is the elementary contraction of $u$ and $v$. It is obvious that both choices yield identical complexes $\mathbb{K}_{i+1}$ up to renaming of vertices. However, the choices make a difference in terms of the size of $\hat{\mathbb{K}}_{i+1}$, because the active closed star of $u$ to $v$ in $\hat{\mathbb{K}}_i$ might differ in size. Our construction simply choose the representative which causes the smaller $\hat{\mathbb{K}}_{i+1}$.

## 3.2 Topological equivalence

We make the following simplifying assumption for $\mathcal{T}$. Let $\phi_i$ be an elementary contraction of $u$ and $v$. If our construction of $\hat{\mathbb{K}}_{i+1}$ turns $v$ inactive, we assume that $\phi_i(u) = \phi_i(v) = u$. Otherwise, we assume $\phi_i(u) = \phi_i(v) = v$. Indeed, this is without loss of generality because it corresponds to a renaming of the simplices in each $\mathbb{K}_i$ and yields equivalent persistence modules. The advantage of this convention is the following property, which follows from a straight-forward inductive argument.

**Lemma 3.** *For every $i$ in $\{0, ..., m\}$, the set of vertices of $\mathbb{K}_i$ is equal to the set of active vertices in $\hat{\mathbb{K}}_i$.*

This allows us to interpret $\mathbb{K}_i$ and $\hat{\mathbb{K}}_i$ as simplicial complexes defined over a common vertex set. In fact, $\mathbb{K}_i$ is the subcomplex of $\hat{\mathbb{K}}_i$ induced by the active vertices:

**Lemma 4.** *A simplex $\sigma$ is in $\mathbb{K}_i$ if and only if $\sigma$ is an active simplex in $\hat{\mathbb{K}}_i$.*

*Proof.* We use induction on $i$. The statement is true for $i = 0$, because $\mathbb{K}_0 = \emptyset = \hat{\mathbb{K}}_0$. So assume first $\phi_i : \mathbb{K}_i \to \mathbb{K}_{i+1}$ is an elementary inclusion that adds a $d$-simplex $\sigma = (v_0, \dots, v_d)$ to $\mathbb{K}_{i+1}$. If $\sigma$ is a vertex, it is set active in $\hat{\mathbb{K}}_{i+1}$ by construction. Otherwise, $v_0, \dots, v_d$ are active in $\hat{\mathbb{K}}_i$ by induction and stay active in $\hat{\mathbb{K}}_{i+1}$. In any case, $\sigma$ is active in $\hat{\mathbb{K}}_{i+1}$. The equivalence for the remaining simplices is straight-forward.

If $\phi_i$ is an elementary contraction $(u, v) \rightsquigarrow u$, we prove both directions of the equivalence separately. For "$\Rightarrow$", fix a $d$-simplex $\sigma \in \mathbb{K}_{i+1}$. It suffices to show that $\sigma \in \hat{\mathbb{K}}_{i+1}$, as in this case, it is also active by Lemma 3. If $\sigma \in \mathbb{K}_i$, this follows at once by induction because $\mathbb{K}_i \subseteq \hat{\mathbb{K}}_i \subseteq \hat{\mathbb{K}}_{i+1}$. If $\sigma \notin \mathbb{K}_i$, $u$ must be a vertex of $\sigma$. Moreover, writing $\sigma = \{u, v_1, ..., v_d\}$ and $\sigma' = \{v, v_1, \dots, v_d\}$ we have that $\sigma' \in \mathbb{K}_i$ and $\phi_i(\sigma') = \sigma$. In particular, the vertices $v_1, \dots, v_d$ are active in $\hat{\mathbb{K}}_i$ by induction, hence $\{v_1, \dots, v_d\}$ is in the active closed star of $v$ in $\hat{\mathbb{K}}_i$. By construction, $\{u, v_1, ..., v_d\} = \sigma$ is in $\hat{\mathbb{K}}_{i+1}$.

For "$\Leftarrow$", let $\sigma \in \hat{\mathbb{K}}_{i+1} \setminus \mathbb{K}_{i+1}$. We show that $\sigma$ is an inactive simplex in $\hat{\mathbb{K}}_{i+1}$. By Lemma 3, this is equivalent to show that $\sigma$ contains a vertex not in $\mathbb{K}_{i+1}$.

Case 1: $\sigma \in \hat{\mathbb{K}}_i$. If $\sigma$ is inactive in $\hat{\mathbb{K}}_i$, it stays inactive in $\hat{\mathbb{K}}_{i+1}$. So, assume that $\sigma$ is active in $\hat{\mathbb{K}}_i$ and thus $\sigma \in \mathbb{K}_i$ by induction. But $\sigma \notin \mathbb{K}_{i+1}$, so $\sigma$ must have $v$ as a vertex and is therefore inactive in $\hat{\mathbb{K}}_{i+1}$.

Case 2: $\sigma \in \hat{\mathbb{K}}_{i+1} \setminus \hat{\mathbb{K}}_i$. By construction, $\sigma$ is of the form $\{u, v_1, \dots, v_d\}$ such that $\{v_1, \dots, v_d\}$ is in the active closed star of $v$ in $\hat{\mathbb{K}}_i$. Assume for a contradiction that $v \neq v_j$ for all $j = 1, \dots, d$. Then, $\sigma' = \{v, v_1, \dots, v_d\}$ is active in $\hat{\mathbb{K}}_i$ and thus, by induction, a simplex in $\mathbb{K}_i$. But then, $\phi_i(\sigma') = \sigma \in \mathbb{K}_{i+1}$ which is a contradiction to our choice of $\sigma$. It follows that $v$ is a vertex of $\sigma$ which proves our claim. □

**Lemma 5.** *For every $0 \leq i \leq m$, the complex $\hat{\mathbb{K}}_i$ collapses to $\mathbb{K}_i$.*

*Proof.* We use induction on $i$. For $i = 0$, $\mathbb{K}_0 = \hat{\mathbb{K}}_0$, and the statement is trivial. Suppose that the statement holds for $\hat{\mathbb{K}}_i$ and $\mathbb{K}_i$ using the sequence $s_i$ of elementary collapses. Note that these collapses only concern inactive simplices in $\hat{\mathbb{K}}_i$. For an inactive vertex $v \in \hat{\mathbb{K}}_i$, the construction of $\hat{\mathbb{K}}_{i+1}$ ensures that $v$ does not gain any additional coface. This implies that $s_i$ is still a sequence of elementary collapses for $\hat{\mathbb{K}}_{i+1}$, yielding a complex $\hat{\mathbb{K}}^*_{i+1}$ with $\mathbb{K}_{i+1} \subseteq \hat{\mathbb{K}}^*_{i+1} \subseteq \hat{\mathbb{K}}_{i+1}$. In particular, $\hat{\mathbb{K}}^*_{i+1}$ only contains vertices that are still active in $\hat{\mathbb{K}}_i$. If $\phi_i$ is an elementary inclusion, $\hat{\mathbb{K}}^*_{i+1} = \mathbb{K}_{i+1}$, because any all vertices in $\hat{\mathbb{K}}_i$ remain active in $\hat{\mathbb{K}}_{i+1}$. For $\phi_i$ being an elementary contraction $(u, v) \rightsquigarrow u$, set $S := \hat{\mathbb{K}}^*_{i+1} \setminus \mathbb{K}_{i+1}$ as the remaining set of simplices that still need to be collapsed to obtain $\mathbb{K}_{i+1}$. All simplices of $S$ have $v$ as vertex. More precisely, $S$ is the set of all simplices of the form $\{v, v_1, \dots, v_d\}$ with $v_1, \dots, v_d$ active in $\hat{\mathbb{K}}_{i+1}$. We split $S = S_u \cup S_{\neg u}$ where $S_u \subset S$ are the simplices in $S$ that contain $u$ as vertex, and $S_{\neg u} = S \setminus S_u$.

We claim that the mapping that sends $\{u, v, v_1, \ldots, v_d\} \in S_u$ to $\{v, v_1, \ldots, v_d\} \in S_{\neg u}$ is a bijection. This map is clearly injective. If $\sigma = \{v, v_1, \ldots, v_d\} \in S_{\neg u}$, then $\sigma \in \hat{\mathbb{K}}_i$ (because every newly added simplex in $\hat{\mathbb{K}}_{i+1}$ contains $u$). Also, $\sigma \in \hat{\mathbb{K}}^*_{i+1}$, and is therefore active in $\hat{\mathbb{K}}_i$. By construction, $\{u, v, v_1, \ldots, v_d\} \in \hat{\mathbb{K}}_{i+1}$, proving surjectivity.

We now define a sequence of elementary collapses from $\hat{\mathbb{K}}^*_{i+1}$ to $\mathbb{K}_{i+1}$. Choose a simplex $\sigma = \{v, v_1, \ldots, v_d\} \in S_{\neg u}$ of maximal dimension, and let $\tau = \{u, v, v_1, \ldots, v_d\}$ denote the corresponding simplex in $S_u$. Then $\sigma$ is indeed a free face in $\hat{\mathbb{K}}^*_{i+1}$, because if there was another coface $\tau' \neq \tau$, it takes the form $\{w, v, v_1, \ldots, v_d\}$ with $w \neq u$ active. So, $\tau' \in S_{\neg u}$, and $\tau'$ has larger dimension than $\sigma$, a contradiction. Therefore, the pair $(\sigma, \tau)$ defines an elementary collapse in $\hat{\mathbb{K}}^*_{i+1}$. We proceed with this construction, always collapsing a remaining pair in $S_{\neg u} \times S_u$ of maximal dimension, until all elements of $S$ have been collapsed. □

**Proposition 6.** *$\mathcal{T}$ and $\mathcal{F}$ have the same barcode.*

*Proof.* Let $\hat{\phi}_i : \hat{\mathbb{K}}_i \to \hat{\mathbb{K}}_{i+1}$ denote the inclusion map from $\hat{\mathbb{K}}_i$ to $\hat{\mathbb{K}}_{i+1}$. By combining Lemma 5 with Lemma 1, we have an isomorphism $\mathrm{inc}_i^* : H(\mathbb{K}_i) \to H(\hat{\mathbb{K}}_i)$, for all $0 \leq i \leq m$, induced by the inclusion maps $\mathrm{inc}_i : \mathbb{K}_i \to \hat{\mathbb{K}}_i$, and therefore the following diagram connecting the persistence modules induced by $\mathcal{T}$ and $\mathcal{F}$:

$$\begin{array}{ccccccc} H(\mathbb{K}_0) & \xrightarrow{\phi_0^*} & H(\mathbb{K}_1) & \xrightarrow{\phi_1^*} & \ldots & \xrightarrow{\phi_{m-1}^*} & H(\mathbb{K}_m) \\ \downarrow{\scriptstyle \mathrm{inc}_0^*} & & \downarrow{\scriptstyle \mathrm{inc}_1^*} & & & & \downarrow{\scriptstyle \mathrm{inc}_m^*} \\ H(\hat{\mathbb{K}}_0) & \xrightarrow{\hat{\phi}_0^*} & H(\hat{\mathbb{K}}_1) & \xrightarrow{\hat{\phi}_1^*} & \ldots & \xrightarrow{\hat{\phi}_{m-1}^*} & H(\hat{\mathbb{K}}_m) \end{array} \tag{1}$$

The *Persistence Equivalence Theorem* [16, p.159] asserts that $(\mathbb{K}_j)_{j=0,\ldots,m}$ and $(\hat{\mathbb{K}}_j)_{j=0,\ldots,m}$ have the same barcode if (1) commutes, that is, if $\mathrm{inc}_{i+1}^* \circ \phi_i^* = \hat{\phi}_i^* \circ \mathrm{inc}_i^*$, for all $0 \leq i < m$.

Two simplicial maps $\phi : \mathbb{K} \to \mathbb{L}$ and $\psi : \mathbb{K} \to \mathbb{L}$ are *contiguous* if, for all $\sigma \in \mathbb{K}$, $\phi(\sigma) \cup \psi(\sigma) \in \mathbb{L}$. Two contiguous maps are known to be homotopic [26, Theorem 12.5.] and thus equal at homology level, that is, $\phi^* = \psi^*$. We show that $\mathrm{inc}_{i+1} \circ \phi_i$ and $\hat{\phi}_i \circ \mathrm{inc}_i$ are contiguous. This implies that (1) commutes, because, by functoriality, $\mathrm{inc}_{i+1}^* \circ \phi_i^* = (\mathrm{inc}_{i+1} \circ \phi_i)^* = (\hat{\phi}_i \circ \mathrm{inc}_i)^* = \hat{\phi}_i^* \circ \mathrm{inc}_i^*$.

To show contiguity, fix $\sigma \in \mathbb{K}_i$ and observe that $(\hat{\phi}_i \circ \mathrm{inc}_i)(\sigma) = \sigma$ because $\hat{\phi}_i$ and $\mathrm{inc}_i$ are inclusions. If $\phi_i(\sigma) = \sigma$, $(\mathrm{inc}_{i+1} \circ \phi_i)(\sigma) = \sigma$ as well, and the contiguity condition is clearly satisfied. So, let $\phi_i(\sigma) \neq \sigma$. Then $\phi_i$ is an elementary contraction $(u, v) \rightsquigarrow u$, and $\sigma$ is of the form $\{v, v_1, \ldots, v_d\}$, where one of the $v_j$ might be equal to $u$. Then, $(\mathrm{inc}_{i+1} \circ \phi_i)(\sigma) = \{u, v_1, ..., v_d\}$. Consequently, $(\mathrm{inc}_{i+1} \circ \phi_i)(\sigma) \cup (\hat{\phi}_i \circ \mathrm{inc}_i)(\sigma) = \{u, v, v_1, ..., v_d\}$. By Lemma 4, $\sigma = \{v, v_1, \ldots, v_d\}$ is in the active closed star of $v$ in $\hat{\mathbb{K}}_i$, and by construction $\{u, v, v_1, ..., v_d\} \in \hat{\mathbb{K}}_{i+1}$, which proves contiguity of the maps. □

## 3.3 Size analysis

**The contracting forest.** We associate a rooted labeled forest $\mathcal{W}_j$ to a prefix $\emptyset = \mathbb{K}_0 \xrightarrow{\phi_0} \ldots \xrightarrow{\phi_{j-1}} \mathbb{K}_j$ of $\mathcal{T}$ inductively as follows: For $j = 0$, $\mathcal{W}_0$ is the empty forest. Let $\mathcal{W}_{j-1}$ be the forest of $\mathbb{K}_0 \to \ldots \to \mathbb{K}_{j-1}$. If $\phi_{j-1}$ is an elementary inclusion of a $d$-simplex, we have two cases: if $d > 0$, set $\mathcal{W}_j := \mathcal{W}_{j-1}$. If a vertex $v$ is included, $\mathcal{W}_j := \mathcal{W}_{j-1} \cup \{x\}$, with $x$ a single node tree labeled with $v$. If $\phi_{j-1}$ is an elementary contraction contracting two vertices $u$ and $v$ in $\mathbb{K}_{j-1}$, there are two trees in $\mathcal{W}_{j-1}$, whose roots are labeled $u$ and $v$. In $\mathcal{W}_j$, these two trees are merged by making their roots children of a new root, which is labeled with the vertex that $u$ and $v$ are mapped to.

We can read off from the construction immediately that the roots of $\mathcal{W}_i$ are labeled with the vertices of the complex $\mathbb{K}_i$, for every $i = 0, \ldots, m$. Moreover, each leaf corresponds to the inclusion of a vertex in $\mathbb{K}_0 \to \ldots \to \mathbb{K}_i$, and each internal node corresponds to a contraction of two vertices. In particular, $\mathcal{W}_i$ is a *full* forest, that is, every node has 0 or 2 children.

Let $\mathcal{W} := \mathcal{W}_m$ denote the forest of the tower $\mathcal{T}$. Let $\Sigma$ denote the set of all simplices that are added at elementary inclusions in $\mathcal{T}$, and recall that $n = |\Sigma|$. A $d$-simplex $\sigma \in \Sigma$ added by $\phi_i$ is formed by $d + 1$ vertices, which correspond to $d + 1$ roots of $\mathcal{W}_{i+1}$, and equivalently, to $d + 1$ nodes

in $\mathcal{W}$. For a node $x$ in $\mathcal{W}$, we denote by $E(x) \subseteq \Sigma$ the subset of simplices with at least one vertex that appears as label in the subtree of $\mathcal{W}$ rooted at $x$. If $y_1$ and $y_2$ are the children of $x$, $E(y_1)$ and $E(y_2)$ are both subsets of $E(x)$, but not disjoint in general. However, the following relation follows at once:

$$|E(x)| \quad \geq \quad |E(y_1)| + |E(y_2) \setminus E(y_1)| \tag{2}$$

We say that the set $N$ of nodes in $\mathcal{W}$ is *independent*, if there are no two nodes $x_1 \neq x_2$ in $N$, such that $x_1$ is an ancestor of $x_2$ in $\mathcal{W}$. A vertex in $\mathbb{K}_i$ appears as label in at most one $\mathcal{W}$-subtree rooted at a vertex in the independent set $N$. Thus, a $d$-simplex $\sigma$ can only appear in up to $d+1$ $E$-sets of vertices in $N$. That implies:

**Lemma 7.** *Let $N$ be an independent set of vertices in $\mathcal{W}$. Then,*

$$\sum_{x \in N} |E(x)| \leq (\Delta + 1) \cdot n,$$

**The cost of contracting.** Recall that a contraction $\mathbb{K}_i \xrightarrow{\phi_i} \mathbb{K}_{i+1}$ yields an inclusion $\hat{\mathbb{K}}_i \hookrightarrow \hat{\mathbb{K}}_{i+1}$ that potentially adds more than one simplex. Therefore, in order to bound the total size of $\hat{\mathbb{K}}_m$, we need to bound the number of simplices added in all these contractions.

We define the *cost* of a contraction $\phi_i$ as $|\hat{\mathbb{K}}_{i+1} \setminus \hat{\mathbb{K}}_i|$, that is, the number of simplices added in this step. Since each contraction corresponds to a node $x$ in $\mathcal{W}$, we can associate these costs to the internal nodes in the forest, denoted by $c(x)$. The leaves get cost 0.

**Lemma 8.** *Let $x$ be an internal node of $\mathcal{W}$ with children $y_1$, $y_2$. Then, $c(x) \leq 2 \cdot |E(y_1) \setminus E(y_2)|$.*

*Proof.* Let $\phi_i : \mathbb{K}_i \to \mathbb{K}_{i+1}$ denote the contraction that is represented by the node $x$, and let $w_1$ and $w_2$ the labels of its children $y_1$ and $y_2$, respectively. By construction, $w_1$ and $w_2$ are vertices in $\mathbb{K}_i$ that are contracted by $\phi_i$. Let $C_1 = \overline{\mathrm{St}}(w_1, \mathbb{K}_i) \setminus \overline{\mathrm{St}}(w_2, \mathbb{K}_i)$ and $C_2 = \overline{\mathrm{St}}(w_2, \mathbb{K}_i) \setminus \overline{\mathrm{St}}(w_1, \mathbb{K}_i)$. By Lemma 4, $\overline{\mathrm{St}}(w_1, \mathbb{K}_i) = \mathrm{Act}\overline{\mathrm{St}}(w_1, \hat{\mathbb{K}}_i)$, and the same for $w_2$. So, because the simplices the two active closed stars have in common will not influence the cost of the contraction, we have,

$$c(x) \leq \min\{|C_1|, |C_2|\}.$$

In particular, $c(x) \leq |C_1|$. We will show that $|C_1| \leq 2 \cdot |E(y_1) \setminus E(y_2)|$.

For every $d$-simplex $\sigma \in \mathbb{K}_i$, there is some $d$-simplex $\tau \in \Sigma$ that has been added in an elementary inclusion $\phi_j$ with $j < i$, such that $\phi_{i-1} \circ \phi_{i-2} \circ \ldots \circ \phi_{j+1}(\tau) = \sigma$. We call $\tau$ an *origin* of $\sigma$. There might be more than one origin of a simplex, but two distinct simplices in $\mathbb{K}_i$ cannot have a common origin. Moreover, for every vertex $v$ of $\sigma$, the tree in $\mathcal{W}_i$ whose root is labeled with $v$ contains exactly one vertex $v'$ of $\tau$ as label. We omit the proof which works by simple induction.

We prove the inequality by a simple charging argument. For each vertex in $C_1$, we charge a simplex in $|E(y_1) \setminus E(y_2)|$ such that no simplex is charged more than twice. Note that $C_1 = (\mathrm{St}(w_1, \mathbb{K}_i) \cup \mathrm{Lk}(w_1, \mathbb{K}_i)) \setminus \overline{\mathrm{St}}(w_2, \mathbb{K}_i)$. If $\sigma \in \mathrm{St}(w_1, \mathbb{K}_i)$, then fix an origin $\tau$ of $\sigma$. Then $\tau$ has a vertex that is a label in the subtree rooted at $y_1$, so $\tau \in E(y_1)$. At the same time, since $w_2$ is not a vertex of $\sigma$, $\tau$ has no vertex in the subtree rooted at $y_2$, so $\tau \notin E(y_2)$. We charge $\tau$ for the existence of $\sigma$. Because different simplices have different origins, every element in $E(y_1) \setminus E(y_2)$ is charged at most once for $\mathrm{St}(w_1, \mathbb{K}_i)$. If $\sigma \in \mathrm{Lk}(w_1, \mathbb{K}_i)$, $\sigma' := w_1 * \sigma \in \mathrm{St}(w_1, \mathbb{K}_i)$, and we can choose an origin $\tau'$ of $\sigma'$. As before, $\tau' \in E(y_1) \setminus E(y_2)$ and we charge $\tau'$ for the existence of $\sigma$. Again, each element in $E(y_1) \setminus E(y_2)$ is charged at most once among all elements in the link. This proves the claim. □

An *ascending path* $(x_1, ..., x_L)$, with $L \geq 1$, is a path in a forest such that $x_{i+1}$ is the parent of $x_i$, for $1 \leq i < L$. We call $L$ the *length* of the path and $x_L$ its *endpoint*. For ascending paths in $\mathcal{W}$, the *cost* of the path is the sum of the costs of the nodes. We say that the set $P$ of ascending paths is *independent*, if the endpoints in $P$ are pairwise different and form an independent set of nodes. We define the *cost* of $P$ as the sum of the costs of the paths in $P$.

**Lemma 9.** *An ascending path with endpoint $x$ has cost at most $2 \cdot |E(x)|$. An independent set of ascending paths in $\mathcal{W}$ has cost at most $2 \cdot (\Delta + 1) \cdot n$.*

*Proof.* For the first statement, let $p = (x_1, ..., x_L)$ be an ascending path with $v_L = v$. Without loss of generality, we can assume the the path starts with a leaf $x_1$, because otherwise, we can always extend the path to a longer path with at least the same cost. We let $p_i = (x_1, \dots, x_i)$ denote the sub-path ending at $x_i$, for $i = 1, \dots, L$, so that $p_L = p$. We let $c(p_i)$ denote the cost of the path $p_i$ and show by induction that $c(p_i) \leq 2 \cdot |E(x_i)|$. For $i = 1$, this follows because $c(p_1) = 0$. For $i = 2, \dots, L$, $x_i$ is an internal node, and its two children are $x_{i-1}$ and some other node $x'_{i-1}$. Using induction and Lemma 8, we have that

$$c(p_i) = c(p_{i-1}) + c(x_i) \leq 2 \cdot (|E(x_{i-1})| + |E(x'_{i-1}) \setminus E(x_{i-1})|) \leq 2 \cdot |E(x_i)|,$$

where the last inequality follows from (2). The second statement follows from Lemma 7 because the endpoints of the paths form an independent set in $\mathcal{W}$. □

**Ascending path decomposition.** An *only-child-path* in a binary tree is an ascending path starting in a leaf and ending at the first encountered node that has a sibling, or at the root of the tree. An only-child-path can have length 1, if the starting leaf has a sibling already. Examples of only-child-paths are shown in Figure 2. We observe that no node with two children lies on any only-child-path, which implies that the set of only-child-paths forms an independent set of ascending paths.

Consider the following pruning procedure for a full binary forest $\mathcal{W}$. Set $\mathcal{W}_{(0)} \leftarrow \mathcal{W}$. In iteration $i$, we obtain the forest $\mathcal{W}_{(i)}$ from $\mathcal{W}_{(i-1)}$ by deleting the only-child-paths of $\mathcal{W}_{(i-1)}$. We stop when $\mathcal{W}_{(i)}$ is empty; this happens eventually because at least the leaves of $\mathcal{W}_{(i-1)}$ are deleted in the $i$-th iteration. Because we start with a full forest, the only-child-paths in the first iteration are all of length 1, and consequently $\mathcal{W}_{(1)}$ arises from $\mathcal{W}_{(0)}$ by deleting the leaves of $\mathcal{W}_{(0)}$. Note that the intermediate forests $\mathcal{W}_{(1)}, \mathcal{W}_{(2)}, \dots$ are not full forests in general. Figure 2 shows the pruning procedure on an example.

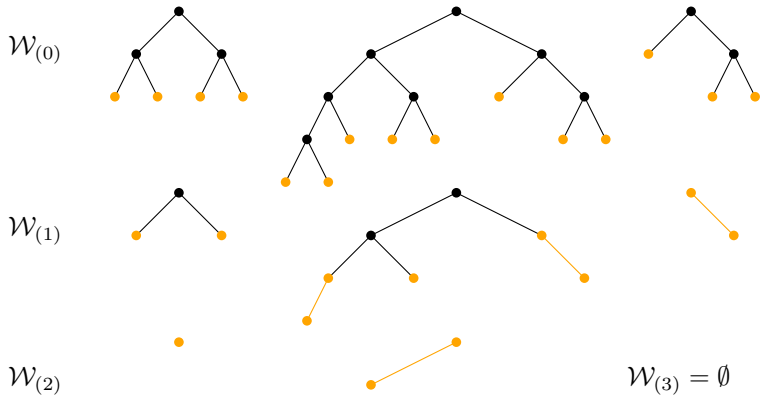


Figure 2: Iterations of the pruning procedure. The only-child-paths are marked in color.

To analyze this pruning procedure in detail, we define the following integer valued function for nodes in $\mathcal{W}$:

$$r(x) = \begin{cases} 1, & \text{if } x \text{ is a leaf} \\ r(y_1) + 1, & \text{if } x \text{ has children } y_1, y_2 \text{ and } r(y_1) = r(y_2) \\ \max\{r(y_1), r(y_2)\}, & \text{if } x \text{ has children } y_1, y_2 \text{ and } r(y_1) \neq r(y_2) \end{cases}$$

**Lemma 10.** *A node $x$ of a full binary forest $\mathcal{W}$ is deleted in the pruning procedure during the $r(x)$-th iteration.*

*Proof.* We prove the claim by induction on the tree structure. If $x$ is a leaf, it is removed in the first iteration, and $r(x) = 1$. If $x$ is an internal node with children $y_1$ and $y_2$, let $r_1 := r(y_1)$ and $r_2 := r(y_2)$. By induction, $y_1$ is deleted in the $r_1$-th iteration and $y_2$ is deleted in the $r_2$-th iteration. There are two cases: if $r_1 = r_2$, $x$ still has two children after $r_1 - 1$ iterations. This implies that $x$ does not lie on an only-child- path in the forest $\mathcal{W}_{(r_1-1)}$ and is therefore not deleted in the $r_1$-th

iteration. But because its children are deleted, $x$ is a leaf in $\mathcal{W}_{(r_1)}$ and therefore deleted in the $(r_1+1)$-th iteration. It remains the second case that $r_1 \neq r_2$. Assume without loss of generality that $r_1 > r_2$, so that $r(x) = r_1$. In iteration $r_1$, $y_1$ lies on an only-child-path in $\mathcal{W}_{(r_1-1)}$. Because $y_2 \notin \mathcal{W}_{(r_1-1)}$, $y_1$ has no sibling in $\mathcal{W}_{(r_1-1)}$, so the only-child-path extends to $x$. Consequently, $x$ is deleted in the $r_1$-th iteration. □

**Lemma 11.** *For a node $x$ in a full binary forest, let $s(x)$ denote the number of nodes in the subtree rooted at $x$. Then $s(x) \geq 2^{r(x)} - 1$. In particular, $r(x) \leq \log_2(s(x)+1)$.*

*Proof.* We prove the claim by induction on $r(x)$. Note that $r(x) = 1$ if and only if $v$ is a leaf, which implies the statement for $r(x) = 1$. For $r(x) > 1$, it is sufficient to prove the statement assuming that $v$ has a minimal $s$-value among all nodes with the same $r$-value. Since $x$ is not a leaf, it has two children $y_1$ and $y_2$. They satisfy $r(y_1) = r(y_2) = r(x) - 1$, because otherwise $r(x) = r(y_1)$ or $r(x) = r(y_2)$, contradicting the minimality of $x$. By induction hypothesis, we obtain that

$$s(x) = 1 + s(y_1) + s(y_2) \geq 1 + \left(2^{r(y_1)} - 1\right) + \left(2^{r(y_2)} - 1\right) = 2^{r(x)} - 1.$$

□

**Proposition 12.** *$|\hat{\mathbb{K}}_m| \leq n + 2 \cdot (\Delta+1) \cdot n \cdot (1 + \log_2(n_0)) = O(n \cdot \Delta \cdot \log_2(n_0))$, where $n_0$ is the number of vertices included in $\mathcal{T}$.*

*Proof.* Applying the pruning procedure to the contraction forest $\mathcal{W}$ of $\mathcal{T}$, we obtain in every iteration a set of independent ascending paths of $\mathcal{W}$, and the cost of this set is bounded by $2 \cdot (\Delta+1) \cdot n$ with Lemma 9. Because $\mathcal{W}$ has at most $2 \cdot n_0 - 1$ nodes, any node $x$ satisfies $r(x) \leq \log_2(2 \cdot n_0)$ by Lemma 11. It follows that the pruning procedure ends after $1 + \log_2(n_0)$ iterations, so the total cost of the contraction forest is at most $2 \cdot (\Delta+1) \cdot n \cdot (1 + \log_2(n_0))$. By definition, this cost is equal to the number of simplices added in all contraction steps. Together with the $n$ simplices added in inclusion steps, the bound follows. □

## 3.4 Algorithm

We will make frequent use of the following concept: a *dictionary* is a data structure that stores a set of *items* of the form `(k,v)`, where `k` is called the *key* and `v` is called the *value* of the item. We assume that all keys stored in the dictionary are pairwise different. The dictionary support three operations: `insert(k,v)` adds a new item to the dictionary, `delete(k)` removes the item with key `k` from the dictionary (if it exists), and `search(k)` returns the item with key `k`, or returns that no such item exists. Common realizations are balanced binary search trees [12, §12] and hash tables [12, §11].

**Simplicial complexes by dictionaries.** The main data structure of our algorithm is a dictionary $D$ that represents a simplicial complex. Every item stored in the dictionary represents a simplex, whose key is the list of its vertices. Every simplex $\sigma$ itself stores an dictionary $CoF_\sigma$. Every item in $CoF_\sigma$ is a pointer to another item in $D$, representing a cofacet $\tau$ of $\sigma$. The key of the item is a vertex identifier (e.g., an integer) for $v$ such that $\tau = v * \sigma$.

How large is $D$ for a simplicial complex $\mathbb{K}$ with $n$ simplices of dimension $\Delta$? Observe that $D$ stores $n$ items, and each key is of size $\leq \Delta+1$. That yields a size of $O(n\Delta)$ plus the size of all $CoF_\sigma$. Since every simplex is the cofacet of at most $\Delta$ simplices, the size of all these inner dictionaries is also bounded by $O(n\Delta)$ (assuming that the size of a dictionary is linear in the number of stored elements).

We can insert and delete simplices efficiently in $D$ using dictionary operations. For instance, to insert a simplex $\sigma$ given as a list of vertices, we insert a new item in $D$ with the key. Then we search for the $\Delta$ facets of $\sigma$ (using dictionary search in $D$), and notify each facet $\tau$ of $\sigma$ about the insertion by adding a pointer to $\sigma$ to $CoF_\tau$, using the vertex $\sigma \setminus \tau$ as the key. The deletion procedure works similarly. Each simplex insertion and deletion requires $O(\Delta)$ dictionary operations. In what follows, it is convenient to assume that dictionary operations have unit costs; we multiply the time complexity with the cost of a dictionary operation at the end to compensate for this simplification.

**The conversion algorithm.** We assume that the tower $\mathcal{T}$ is given to us as a stream where each element represents a simplicial map $\phi_i$ in the tower. Specifically, an element starts with a token {`INCLUSION`, `CONTRACTION`} that specifies the type of the map. In the first case, the token is followed by a non-empty list of vertex identifiers specifying the vertices of the simplex to be added. In the second case, the token is followed by two vertex identifiers $u$ and $v$, specifying a contraction of type $(u, v) \rightsquigarrow u$.

The algorithm outputs a stream of simplices specifying the filtration $\mathcal{F}$. Specifically, while handling the $i$-th input element, it outputs the simplices of $\hat{\mathbb{K}}_{i+1} \setminus \hat{\mathbb{K}}_i$ in increasing order of dimension (to ensure that every prefix is a simplicial complex). For simplicity, we assume that output simplices are also specified by a list of vertices – the algorithm can easily be adapted to return the boundary matrix of the filtration in sparse list representation with the same complexity bounds.

We use an initially empty dictionary $D$ as above, and maintain the invariant that after the $i$-th iteration, $D$ represents the active subcomplex of $\hat{\mathbb{K}}_i$, which is equal to $\mathbb{K}_i$ by Lemma 4.

If the algorithm reads an inclusion of a simplex $\sigma$ from the stream, it simply adds $\sigma$ to $D$ (maintaining the invariant) and writes $\sigma$ to the output stream.

If the algorithm reads a contraction of two vertices $u$ and $v$, from $\mathbb{K}_i$ to $\mathbb{K}_{i+1}$, we let $c_i$ denote the cost of the contraction, that is, $c_i = |\hat{\mathbb{K}}_{i+1} \setminus \hat{\mathbb{K}}_i|$. The first step is to determine which of the vertices has the smaller active closed star in $\hat{\mathbb{K}}_i$, or equivalently, which vertex has the smaller closed star in $\mathbb{K}_i$. The size of the closed star of a vertex $v$ could be computed by a simple graph traversal in $D$, starting at a vertex $v$ and following the cofacet pointers recursively, counting the number of simplices encountered. However, we want to identify the smaller star with only $O(c_i)$ operations, and the closed star can be much larger. Therefore, we change the traversal in several ways:

First of all, observe that $|\overline{\mathrm{St}}(u)| \leq |\overline{\mathrm{St}}(v)|$ if and only if $|\mathrm{St}(u)| \leq |\mathrm{St}(v)|$ (in $\mathbb{K}_i$). Now define $\mathrm{St}(u, \neg v) := \mathrm{St}(u) \setminus \mathrm{St}(v)$. Then, $|\mathrm{St}(u)| \leq |\mathrm{St}(v)|$ if and only if $|\mathrm{St}(u, \neg v)| \leq |\mathrm{St}(v, \neg u)|$, because we subtracted the intersection of the stars on both sides. Finally, note that $\min\{|\mathrm{St}(u, \neg v)|, |\mathrm{St}(v, \neg u)|\} \leq c_i$, as one can easily verify. Moreover, we can count the size of $\mathrm{St}(u, \neg v)$ by a cofacet traversal from $u$, ignoring cofacets that contain $v$ (using the key of $CoF_*$), in $O(|\mathrm{St}(u, \neg v)|)$ time. However, this is still not enough, because counting both sets independently gives a running time of $\max\{|\mathrm{St}(u, \neg v)|, |\mathrm{St}(v, \neg u)|\}$. The last trick is that we count the sizes of $\mathrm{St}(u, \neg v)$ and $\mathrm{St}(v, \neg u)$ at the same time by a simultaneous graph traversal of both, terminating as soon as one of the traversal stops. The running time is then proportional to $2 \cdot \min\{|\mathrm{St}(u, \neg v)|, |\mathrm{St}(v, \neg u)|\} = O(c_i)$, as required.

Assume w.l.o.g. that $|\overline{\mathrm{St}}(u)| \leq |\overline{\mathrm{St}}(v)|$. Also in time $O(c_i)$, we can obtain $\mathrm{St}(u, \neg v)$. We sort its elements by increasing dimension, which can be done in $O(c_i + \Delta)$ using integer sort. For each simplex $\sigma = \{u, v_1, \ldots, v_k\} \in \mathrm{St}(u, \neg v)$ in order, we check whether $\{v, v_1, \ldots, v_k\}$ is in $D$. If not, we add it to $D$ and also write it to the output stream. Then, we write $\{u, v, v_1, \ldots, v_k\}$ to the output stream (note that we do not have to check it existence in $\hat{\mathbb{K}}_i$, because it does not by construction, and there is no need to add it to $D$ because of the next step). At the end of the loop, we wrote exactly the simplices in $\mathbb{K}_{i+1} \setminus \mathbb{K}_i$ to the output stream, which proves correctness.

It remains to maintain the invariant on $D$. Assuming still that $|\overline{\mathrm{St}}(u)| \leq |\overline{\mathrm{St}}(v)|$, $u$ turns inactive in $\hat{\mathbb{K}}_{i+1}$. We simply traverse over all cofaces of $u$ and remove all encountered simplices from $D$. After this operations, the invariant holds. This ends the description of the algorithm.

**Complexity analysis.** The following statement combined with Propositions 6 and 12 completes the proof of Theorem 2.

**Proposition 13.** *The algorithm requires $O(\Delta \cdot \omega)$ space and $O(\Delta \cdot |\hat{\mathbb{K}}_m| \cdot C_\omega)$ time, where $\omega = \max_{i=0,\ldots,m} |\mathbb{K}_i|$ and $C_\omega$ is the cost of an operation in a dictionary with at most $\omega$ elements.*

*Proof.* The space complexity follows at once from the invariant, because the size of $D$ is at most $O(\Delta|\mathbb{K}_i|)$ during the $i$-th iteration.

For the time bound, we set $S := |\hat{\mathbb{K}}_m|$ for convenience and show that the algorithm finishes in $O(\Delta \cdot S)$ steps, assuming dictionary operations to be of constant cost. We have one simplex insertion per elementary inclusion which requires $O(\Delta)$ operations. Thus, all inclusions are bounded by $O(n\Delta)$, which is subsumed by our bound as $n \leq S$. For the contraction case, we need $O(c_i)$ to identify the smaller star, $O(c_i + \Delta)$ to get a sorted list of $\mathrm{St}(u, \neg v)$ (or vice versa), and $O(\Delta \cdot c_i)$ to

add new vertices to $D$. Moreover, we delete the star of $u$ from $D$; the cost for that is $O(\Delta \cdot d_i)$, where $d_i$ is the number of deleted simplices. Thus, the complexity of a contraction is $O(\Delta \cdot (c_i + d_i))$.

Since $c_i$ is the number of simplices added to the filtration at step $i$, the sum of all $c_i$ is bounded by $O(S)$. Moreover, because every simplex that ever appears in $D$ belongs to $\hat{\mathbb{K}}_m$ and every simplex is inserted only once, the sum of all $d_i$ is bounded by $O(S)$ as well. Therefore, the combined cost over all contractions is $O(\Delta \cdot S)$ as required. □

Note that the dictionary $D$ has lists of identifiers of length up to $\Delta+1$ as keys, so that comparing two keys has cost $O(\Delta)$. Therefore, using balanced binary trees as dictionary, we get a complexity of $C_\omega = O(\Delta \log \omega)$. Using hash tables, we get an expected worst-case complexity of $C_\omega = O(\Delta)$.

**Experimental results.** The following tests where made on a 64-bit Linux (Ubuntu) HP machine with a 3.50 GHz Intel processor and 63 GB RAM. The programs were all implemented in C++ and compiled with optimization level -O2.

To test the performance of our algorithm, we compared it to the software Simpers, which is the implementation of the Annotation Algorithm from Dey, Fan and Wang described in [14]. Simpers computes the persistence of the given filtration, so we add to our time the time the library PHAT needs to compute the persistence from the filtration we generate. We used PHAT with the default parameters ('–twist' and '–bit_tree_pivot_column').

The results of the tests are in Table 1. The timings for File IO are not included in the process time of PHAT and Simpers. The symbol $\infty$ means that the calculation time exceeded 12 hours. The memory peak was obtained via the '/usr/bin/time -v' Linux command. We used in this table two types of data: the first three filtrations, `data1-3`, were generated randomly, where the second ones are Batch-collapsed Rips filtration computed from particular point clouds, taken from [15].

| | | | | | | Alg1 + PHAT | | | Simpers | |
|---|---|---|---|---|---|---|---|---|---|---|
| | $c$ | $n$ | $n_0$ | $\Delta$ | $\omega$ | filtration size | time (s) | mem. peak (kB) Alg1 / total | time (s) | mem. peak (kB) |
| data1 | 495 | 4 833 | 500 | 4 | 2 908 | 19 747 | 0.12 | 4 644 / 7 040 | 2.49 | 10188 |
| data2 | 795 | 7 978 | 800 | 4 | 4 816 | 35 253 | 0.20 | 5 424 / 10 228 | 13.97 | 20308 |
| data3 | 794 | 8 443 | 800 | 5 | 5 155 | 38 101 | 0.22 | 5 744 / 10 916 | 19.29 | 24 924 |
| GPS | 1 746 | 8 585 | 1 747 | 3 | 1 747 | 9 063 | 0.07 | 4 072 / 5 292 | 0.35 | 6 064 |
| KB | 22 499 | 95 019 | 22 500 | 3 | 22 500 | 133 433 | 0.50 | 10 520 / 18 712 | 2.83 | 24 460 |
| MC | 23 074 | 143 928 | 23 075 | 3 | 28 219 | 185 447 | 0.72 | 14 636 / 25 272 | 4.12 | 26 020 |
| S3 | 252 995 | 1 473 580 | 252 996 | 4 | 252 996 | 1 824 461 | 10.09 | 94 020 / 221 636 | 49.86 | 239 404 |
| PC25 | 14 999 | 10 246 125 | 15 000 | 3 | 2 191 701 | 12 283 003 | 135.02 | 1 029 680 / 2 223 544 | $\infty$ | - |

Table 1: Experimental results.

To verify that the space consumption of our algorithm does not dependent on the length of the tower, we constructed an additional example whose size exceed our RAM capacity, but whose width is small: we took 10 random points moving on a flat torus in a randomly chosen fixed direction. When two points get in distance less than $t_1$ to each other, we add the edge between them (the edge remains also if the points increase their distance later on). When two points get in distance less than $t_2$ from each other with $t_2 < t_1$, we contract the two vertices and let a new moving point appear somewhere on the torus. This process yields a sequence of graphs, and we take its flag complex as our simplicial tower. In this way, we obtain a tower of length about $3.5 \cdot 10^9$ which has a file size of about 73 GB, but only has a width of 367. Our algorithm took about 2 hours to convert this tower into a filtration of size roughly $4.6 \cdot 10^9$. During the conversion, the virtual memory used was constantly around 22 MB and the resident set size about 3.8 MB only, confirming the theoretical prediction that the space consumption is independent of the length of the tower. The information about the memory use was taken from the '/proc/<pid>/stat' system file every 100 000 insertions/contractions during the process.

# 4 Persistence by Streaming

Even if the complex we need to maintain in memory during the conversion is relatively small, at the end, the algorithm to compute the final persistence still needs to memorize the whole filtration we give it as input. If the complex in the original tower has an interesting maximum size during the whole process, we should be able to compute its persistence even if the tower is extremely long. So we design here a streaming variation of the reduction algorithm that computes the barcode of filtrations, such that it has an efficient memory use. More precisely, we will prove the following theorem:

**Theorem 14.** *With the same notation as in Theorem 14, we can compute the barcode of a tower $\mathcal{T}$ in worst-case time $O(\omega^2 \cdot \Delta \cdot n \cdot \log n_0)$ and space complexity $O(\omega^2)$*

We describe the algorithm in Section 4.1 and prove the complexity bounds in Section 4.2. The described algorithm requires various adaptations to become efficient in practice, and we describe an improved variant in Section 4.3. We present some experiments in Section 4.4.

## 4.1 Algorithmic description

On a high level, our algorithm converts the tower into an equivalent filtration and computes the barcode of that filtration. We focus on the second part, which we describe as a streaming algorithm. The input to the algorithm is a stream of elements, each starting with a token {`ADDITION`, `INACTIVE`} followed by a simplex identifier which represents a simplex $\sigma$. In the addition case, this is followed by a list of simplex identifiers specifying the facets of $\sigma$. In other words, the element encodes the next column of the boundary matrix. For the inactive case, it means that $\sigma$ has become inactive in the complex. In particular, no subsequent simplex in the stream will contain $\sigma$ as a facet. It is not difficult to modify the algorithm from Section 3.4 to return a stream as required, within the same complexity bounds.

The algorithm uses a matrix data type $M$ as its main data structure. We realize $M$ as a dictionary of columns, indexed by a simplex identifier. Each column is a sorted linked list of identifiers corresponding to the non-zero row indices of the column. In particular, we can access the pivot of the column in constant time and we can add two columns in time proportional to the maximal size of the involved lists. Note that most algorithms store the boundary matrix as an array of columns, but we use dictionaries for space efficiency.

There are two secondary data structures that we mention briefly: given a row index $r$, we have to identify the column index $c$ of the column that has $r$ as pivot in the matrix (or to find out that no such column exists). This can be done using a dictionary with key and value type both simplex identifiers. Finally, we maintain a dictionary representing the set of simplex identifiers that represent active simplices of the complex, plus a flag denoting whether the corresponding simplex is positive or negative. It is straight-forward to maintain these structures during the algorithm, and we will omit the description of the required operations.

The algorithm uses two subroutines. The first one, called `reduce_column`, takes a column identifier $j$ as input is defined as follows: iterate through the non-zero row indices of $j$. If an index $i$ is the index of an inactive and negative column in $M$, remove the entry from the column $j$ (this is the "compression" described at the end of Section 2). After this pre-processing, reduce the column: while the column is non-empty, and its pivot $i$ is the pivot of another column $k < j$ in the matrix, add column $k$ to column $j$.

The second subroutine, `remove_row`, takes a index $\ell$ as input and clears out all entries in row $\ell$ from the matrix. For that, let $j$ be the column with $\ell$ as pivot. Traverse all non-zero columns of the matrix except column $j$. If a column $i \neq j$ has a non-zero entry at row $\ell$, add column $j$ to column $i$. After traversing all columns, remove column $j$ from $M$.

The main algorithm can be described easily now: if the input stream contains an addition of a simplex, we add the column to $M$ and call `reduce_column` on it. If at the end of that routine, the column is empty, it is removed from $M$. If the column is not empty and has pivot $\ell$, we report $(\ell, j)$ as a persistence pair and check whether $\ell$ is active. If not, we call `remove_row(`$\ell$`)`. If the input stream specifies that simplex $\ell$ becomes inactive, we check whether $j$ appears as pivot in the matrix and call `remove_row(`$\ell$`)` in this case.

**Proposition 15.** *The algorithm computes the correct barcode.*

*Proof.* First, note that removing a column from $M$ within the procedure `remove_row` does not affect further reduction steps: Since the pivot $\ell$ of the column is inactive, no subsequent column in the stream will have an entry in row $\ell$. Moreover, the reduction process cannot introduce an entry in row $\ell$ because the routine has removed all such entries.

Note that `remove_row` might also include right-to-left column additions, and we also have to argue that they do not change the pivots. Let $R$ denote the matrix obtained by the standard compression algorithm that does not call `remove_row` (as described in Section 2). Just before our algorithm calls `reduce_column` on the $j$-th column, let $S$ denote the matrix with $(j-1)$ columns that is represented by $M$. It is straight-forward to verify by an inductive argument that every column of $S$ is a linear combination of $R_1, \dots, R_{j-1}$. `reduce_column` adds a subset of the columns of $S$ to the $j$-th column. Thus, the reduced column can be expressed by a sequence of left-to-right column additions in $R$, and thus yields the same pivot as the standard compression algorithm. □

## 4.2 Complexity analysis

We analyze how large the structure $M$ can become during the algorithm. After every iteration, the matrix represents the reduced boundary matrix of some intermediate complex $\hat{\mathbb{L}}$ with $\hat{\mathbb{K}}_i \subseteq \hat{\mathbb{L}} \subseteq \hat{\mathbb{K}}_{i+1}$ for some $i = 0, \dots, m$. Moreover, the active simplices define a subcomplex $\mathbb{L} \subseteq \hat{\mathbb{L}}$ and there is a moment during the algorithm where $\hat{\mathbb{L}} = \hat{\mathbb{K}}_i$ and $\mathbb{L} = \mathbb{K}_i$, for every $i = 0, \dots, m$. We call this the *$i$-th checkpoint*. We will make frequent use of the following simple observation.

**Lemma 16.** $|\hat{\mathbb{K}}_{i+1} \setminus \hat{\mathbb{K}}_i| \leq |\mathbb{K}_i| \leq \omega$

**Lemma 17.** *At every moment, the number of columns stored in $M$ is at most $2\omega$.*

*Proof.* It can be verified easily that throughout the algorithm, a column is stored in $M$ only if not zero, and its pivot is active. So, assume first that we are at the $i$-th checkpoint for some $i$. Since $\mathbb{L} = \mathbb{K}_i$, there are not more than $|\mathbb{K}_i| \leq \omega$ active simplices. Since each column has a different active pivot, their number is also bounded by $\omega$. If we are between checkpoint $i$ and $i+1$, there have been not more than $\omega$ columns added to $M$ since the $i$-th checkpoint from Lemma 16. The bound follows. □

The number of rows is more difficult to bound because we cannot guarantee that each column in $M$ corresponds to an active simplex. Still, the number of rows is asymptotically the same as for columns:

**Lemma 18.** *At every moment, the number of rows stored in $M$ is at most $4\omega$.*

*Proof.* Consider a row index $\ell$ and a time in the algorithm where $M$ represents $\hat{\mathbb{L}}$. By the same argument as in the previous lemma, we can argue that there are at most $2\omega$ active row indices at any time. Therefore, we restrict our attention to the case that $\ell$ is inactive and distinguish three cases. If $\ell$ represents a negative simplex, we observe that its row should have been was removed due to the compression optimization, and after $\ell$ became inactive, no simplex can have it as a facet either. It follows that row $\ell$ is empty in this case. If $\ell$ is positive and was paired with another index $j$ during the algorithm, then `remove_row` was called on $\ell$, either at the moment the pair was formed, or when $\ell$ became inactive. Since the procedure removes the row, we can conclude that row $\ell$ is empty also in this case. The final case is that $\ell$ is positive, but has not been paired so far. It is well-known that in this case, $\ell$ is the generator of an homology class of $\hat{\mathbb{L}}$. Let

$$\beta(\hat{\mathbb{L}}) := \sum_{i=0}^{\Delta} \beta_i(\hat{\mathbb{L}})$$

denote the sum of the Betti numbers of the complex. Then, it follows that the number of such row indices is at most $\beta(\hat{\mathbb{L}})$.

We argue that $\beta(\hat{\mathbb{L}}) \leq 2\omega$ which proves our claim. Assume that $\hat{\mathbb{K}}_i \subseteq \hat{\mathbb{L}} \subset \hat{\mathbb{K}}_{i+1}$. We have that $\beta(\hat{\mathbb{K}}_i) = \beta(\mathbb{K}_i)$ by Lemma 5, and since $\mathbb{K}_i$ has at most $\omega$ simplices, $\beta(\mathbb{K}_i) \leq \omega$. Since we add at most $\omega$ simplices to get from $\hat{\mathbb{K}}_i$ to $\hat{\mathbb{L}}$, and each addition can increase $\beta$ by at most one, we have indeed that $\beta(\hat{\mathbb{L}}) \leq 2\omega$. □

**Proposition 19.** *The algorithm has time complexity* $O(\omega^2 \cdot \Delta \cdot n \cdot \log n_0)$ *and space complexity* $O(\omega^2)$.

*Proof.* The space complexity is immediately clear from the preceding two lemmas, as $M$ is the dominant data structure in terms of space consumption. For the time complexity, we observe that both subroutines `reduce_column` and `remove_row` need $O(\omega)$ column additions and $O(\omega)$ dictionary operations in the worst case. A column addition costs $O(\omega)$, and a dictionary operation is not more expensive (since the dictionaries contain at most $O(\omega)$ elements and their keys are integers). So, the complexity of both methods is $O(\omega^2)$. Since each routine is called at most once per input element, and there are $O(\Delta \cdot n \cdot \log n_0)$ elements by Theorem 2, the bound follows. □

## 4.3 Implementation

The algorithm described in Section 4.1 is not efficient in practice for three reasons: First, it has been observed as a general rule that the *clearing optimization* [3] significantly improves the standard algorithm. However, in the above form, that optimization is not usable because it requires to process the columns in a non-incremental way. Second, the `remove_row` routine scans the entire matrix $M$; while not affecting the worst-case complexity, frequently scanning the matrix should be avoided in practice. Finally, the above algorithm uses lists to represent columns, but it has been observed that this is a rather inefficient way to perform matrix operations [3].

We outline a variant of the above algorithm that partially overcomes these drawbacks and behaves better in practice. In particular, our variant can be implemented with all column representations available in the PHAT library. The idea is to perform a "batch" variant of the previous algorithm: We define a *chunk size* $C$ and read in $C$ elements from the stream; we insert added columns in the matrix, removing row entries of already known inactive negative columns as before, but not reducing the columns yet. After having read $C$ elements, we start the reduction of the newly inserted columns using the clearing optimization. That is, we go in decreasing dimension and remove a column as soon as its index becomes the pivot of another column; see [9] for details. After the reduction ends, except for the last chunk, we go over the columns of the matrix and check for each pivot whether it is active. If it is, we traverse its row entries in decreasing order, but skipping the pivot. Let $\ell$ be the current entry. If $\ell$ is the inactive pivot of some column $j$, we add $j$ to the current column. If $\ell$ is inactive and represents a negative column, we delete $\ell$ from the current column. After performing these steps for all remaining columns of the matrix, we go over all columns again, deleting every column with inactive pivot. (By the way, also cleaning up the secondary data structures described in 4.1.)

It remains the question of how to choose the parameter $C$. The chunk provides a trade-off between time and space efficiency. Roughly speaking, the matrix can have up to $O(\omega + C)$ columns during this reduction, but the larger the chunks are, the more benefit one can draw from the clearing optimization (the clearing optimization fails for pairs where the simplices are in different chunks). We therefore recommend to choose $C$ rather large, but making sure that the matrix will still fit into memory.

## 4.4 Experimental evaluation

The tests were made with the same setup as in Section 3.4. Figure 3 shows the effect of the chunk size parameter $C$ on the runtime and memory consumption of the algorithm. The data used is `S3` (see Section 3.4); we also performed the tests on the other examples from Table 1, with similar outcome. The File IO operations are included in the measurements. Confirming the theory, as the chunk size decreases, our implementation needs less space but more computation time (while the running time seems to increase slightly again for larger chunk sizes).

For the $4.6 \cdot 10^9$ inclusions tower from Section 3.4, with $C = 200\,000$, the algorithm took around 4.5 hours, the virtual memory used was constantly around 68 MB and the resident set size constantly around 49 MB, confirming the theoretical statement that the memory size does not depend on the length of the filtration.

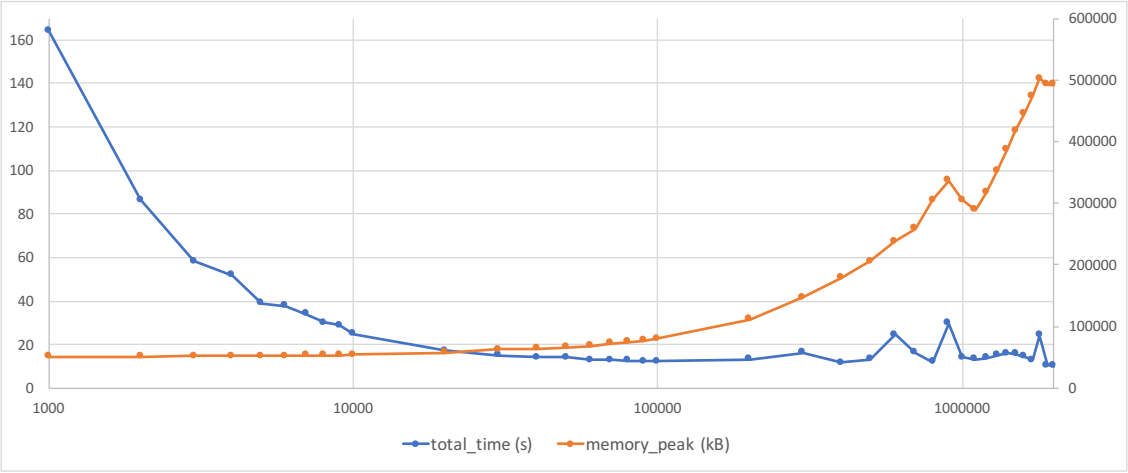


Figure 3: Evolution of processing time (left Y-axis in sec) and process memory peak (right Y-axis in kB) depending on the chunk size (logarithmic X-axis).

# References


[1] U. Bauer, M. Kerber, and J. Reininghaus. Clear and compress: Computing persistent homology in chunks. In *Topological Methods in Data Analysis and Visualization III*, Mathematics and Visualization, pages 103–117. Springer, 2014.

[2] U. Bauer, M. Kerber, and J. Reininghaus. Distributed computation of persistent homology. In *Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 31–38, 2014.

[3] U. Bauer, M. Kerber, J. Reininghaus, and H. Wagner. Phat – persistent homology algorithms toolbox. *Journal of Symbolic Computation*, 78:76–90, 2017.

[4] J.-D. Boissonnat, T. Dey, and C. Maria. The compressed annotation matrix: An efficient data structure for computing persistent cohomology. In *European Symposium on Algorithms (ESA)*, pages 695–706, 2013.

[5] M. Botnan and G. Spreemann. Approximating Persistent Homology in Euclidean space through collapses. *Applied Algebra in Engineering, Communication and Computing*, 26(1-2):73–101, 2015.

[6] G. Carlsson. Topology and data. *Bulletin of the AMS*, 46:255–308, 2009.

[7] G. Carlsson and V. de Silva. Zigzag persistence. *Foundations of Computational Mathematics*, 10(4):367–405, 2010.

[8] G. Carlsson, V. de Silva, and D. Morozov. Zigzag persistent homology and real-valued functions. In *ACM Symposium on Computational Geometry (SoCG)*, pages 247–256, New York, NY, USA, 2009. ACM.

[9] C. Chen and M. Kerber. Persistent homology computation with a twist. In *European Workshop on Computational Geometry (EuroCG)*, pages 197–200, 2011.

[10] C. Chen and M. Kerber. An output-sensitive algorithm for persistent homology. *Computational Geometry: Theory and Applications*, 46(4):435–447, 2013.

[11] A. Choudhary, M. Kerber, and S. Raghvendra. Polynomial-sized topological approximations using the permutahedron. In *32nd International Symposium on Computational Geometry (SoCG 2016)*, pages 31:1–31:16.

[12] T. Cormen, C. Leiserson, R. Rivest, and C. Stein. *Introduction to algorithms*. The MIT press, 3rd edition, 2009.

[13] V. de Silva, D. Morozov, and M. Vejdemo-Johansson. Dualities in persistent (co)homology. *Inverse Problems*, 27(12):124003+, 2011.

[14] T. Dey, F. Fan, and Y. Wang. Computing Topological Persistence for Simplicial Maps. In *ACM Symposium on Computational Geometry (SoCG)*, pages 345–354, 2014.

[15] T. Dey, D. Shi, and Y. Wang. Simba: An efficient tool for approximating rips-filtration persistence via simplicial batch-collapse. In *European Symposium on Algorithms (ESA)*, 2016. to appear.

[16] H. Edelsbrunner and J. Harer. *Computational Topology: an introduction*. American Mathematical Society, 2010.

[17] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28(4):511–533, 2002.

[18] H. Edelsbrunner and S. Parsa. On the computational complexity of betti numbers: Reductions from matrix rank. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 152–160, 2014.

[19] P. Gabriel. Unzerlegbare Darstellungen I. *manuscripta mathematica*, 6(1):71–103, 1972.

[20] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.

[21] M. Kerber. Persistent homology: State of the art and challenges. *Internationale Mathematische Nachrichten*, 231:15–33, 2016.

[22] M. Kerber and R. Sharathkumar. Approximate Čech complex in low and high dimensions. In *International Symposium on Algortihms and Computation (ISAAC)*, pages 666–676, 2013.

[23] C. Maria, J.-D. Boissonnat, M. Glisse, and M. Yvinec. The Gudhi library: Simplicial complexes and persistent homology. In *International Congress on Mathematical Software (ICMS)*, volume 8592 of *Lecture Notes in Computer Science*, pages 167–174, 2014.

[24] C. Maria and S. Oudot. Zigzag persistence via reflections and transpositions. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, pages 181–199. Society for Industrial and Applied Mathematics, 2015.

[25] N. Milosavljevic, D. Morozov, and P. Skraba. Zigzag persistent homology in matrix multiplication time. In *ACM Symposium on Computational Geometry (SoCG)*, pages 216–225, 2011.

[26] J. Munkres. *Elements of Algebraic Topology*. Perseus Publishing, 1984.

[27] N. Otter, M. Porter, U. Tillmann, P. Grindrod, and H. Harrington. A roadmap for the computation of persistent homology. *arXiv*, abs/1506.08903, 2015.

[28] S. Oudot. *Persistence theory: From Quiver Representation to Data Analysis*, volume 209 of *Mathematical Surveys and Monographs*. American Mathematical Society, 2015.

[29] D. Sheehy. Linear-size approximation to the Vietoris-Rips filtration. *Discrete & Computational Geometry*, 49:778–796, 2013.

[30] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.